

CLAIMS

1. An asynchronous completion notification system for use in an RDMA (remote data memory access) network interface card (RNIC) having a completion queue (CQ) for holding completion queue entries (CQEs), comprising:

a system for storing a first CQE number of the most recent CQE placed into the CQ;

a system for storing a second CQE number of the most recent CQE retrieved from the CQ;

a system for packaging the second CQE number with each request completion notification verb that is issued; and

a processing system for processing the request completion notification verb, wherein the processing system compares the first CQE number with the second CQE number to determine whether asynchronous completion notification should be immediately performed.

2. The asynchronous completion notification system of claim 1, wherein the first CQE number is stored in a CQ context.

3. The asynchronous completion notification system of claim 1, wherein the system for storing the second CQE number is implemented by a verb layer.

4. The asynchronous completion notification system of claim 3, wherein the verb layer packages the second CQE number with the request completion notification verb.
5. The asynchronous completion notification system of claim 1, wherein the processing system causes asynchronous completion notification to be immediately performed if the second CQE number is less than the first CQE number.
6. The asynchronous completion notification system of claim 1, wherein the processing system waits until a next CQE is placed into the CQ before performing asynchronous completion notification if the second CQE number is equal to the first CQE number.
7. The asynchronous completion notification system of claim 1, further comprising:
 - a system for storing a third CQE number for the most recent solicited CQE placed into the CQ,
 - wherein the request completion notification verb is further packaged with a type of completion notification;
 - wherein the processing system includes a system for checking the type of completion notification; and
 - wherein, if the type of completion notification is solicited, the processing compares the third CQE number with the second CQE number to determine whether asynchronous completion notification should be immediately performed.

8. The asynchronous completion notification system of claim 7, wherein if the type of completion notification is solicited, the processing system causes asynchronous completion notification to be immediately performed if the second CQE number is less than the third CQE number.

9. The asynchronous completion notification system of claim 7, wherein if the type of completion notification is solicited, the processing system waits until a next solicited CQE is placed into the CQ before causing asynchronous completion notification to be performed if the second CQE number is equal to or greater than the third CQE number.

10. A method for implementing asynchronous completion notification in an RDMA (remote data memory access) network interface card (RNIC) having a completion queue (CQ) for holding completion queue entries (CQEs), comprising:

- storing a first CQE number of a most recent CQE placed into the CQ;
- storing a second CQE number of a most recent CQE retrieved from the CQ;
- issuing a request for completion notification;
- packaging the second CQE number with the request; and
- processing the request, wherein the processing step compares the first CQE number with the second CQE number to determine whether asynchronous completion notification should be immediately performed.

11. The method of claim 10, wherein the first CQE number is stored in a CQ context.

12. The method of claim 10, wherein a verb layer packages the second CQE number with a request completion notification verb.

13. The method of claim 10, wherein the processing step immediately performs asynchronous completion notification if the second CQE number is less than the first CQE number.

14. The method of claim 10, wherein the processing step waits until a next CQE is placed into the CQ before performing asynchronous completion notification if the second CQE number is equal to the first CQE number.

15. The method of claim 10, further comprising:

storing a third CQE number for the most recent solicited CQE placed into the CQ;
packaging a type of completion notification with the request;
checking the type of completion notification during the processing step; and
if the type of completion notification is solicited, comparing the third CQE number with the second CQE number to determine whether asynchronous completion notification should be immediately performed.

16. The method of claim 15, wherein if the type of completion notification is solicited, the processing step immediately performs asynchronous completion notification if the second CQE number is less than the third CQE number.

17. The method of claim 16, wherein if the type of completion notification is solicited, the processing step waits until a next solicited CQE is placed into the CQ before performing asynchronous completion notification if the second CQE number is equal to or greater than the third CQE number.

18. A system for implementing asynchronous completion notification in an RDMA (remote data memory access) network interface card (RNIC) having a completion queue (CQ) for holding completion queue entries (CQEs), comprising:

means for storing a first CQE number of a most recent CQE placed into the CQ;

means for storing a second CQE number of a most recent CQE retrieved from the CQ;

means for issuing a request for completion notification;

means for packaging the second CQE number with the request; and

means for processing the request, wherein the processing means compares the first CQE number with the second CQE number to determine whether asynchronous completion notification should be immediately performed.

19. The system of claim 18, wherein the processing means immediately performs asynchronous completion notification if the second CQE number is less than the first CQE number.

20. The system of claim 18, wherein the processing means waits until a next CQE is placed into the CQ before performing asynchronous completion notification if the second CQE number is equal to the first CQE number.

21. The system of claim 18, further comprising:

means for storing a third CQE number for the most recent solicited CQE placed into the CQ;

means for packaging a type of completion notification with the request;

means for checking the type of completion notification; and

means for comparing the third CQE number with the second CQE number if the type of completion notification is solicited, to determine whether asynchronous completion notification should be immediately performed.

22. The system of claim 21, wherein the processing means immediately performs asynchronous completion notification if the type of completion notification is solicited and if the second CQE number is less than the third CQE number.

23. The system of claim 21, wherein the processing means waits until a next solicited CQE is placed into the CQ before performing asynchronous completion notification if the type of completion notification is solicited and if the second CQE number is equal to or greater than the third CQE number.